
Random Sampling in an Age of Automation: Minimizing Expenditures through Balanced Collection and Annotation

Oscar Beijbom
UC Berkeley
obeijbom@berkeley.edu

Abstract

Methods for automated collection and annotation are changing the cost-structures of sampling surveys for a wide range of applications. Digital samples in the form of images or audio recordings can be collected rapidly, and annotated by computer programs or crowd workers. We consider the problem of estimating a population mean under these new cost-structures, and propose a Hybrid-Offset sampling design. This design utilizes two annotators: a primary, which is accurate but costly (e.g. a human expert) and an auxiliary which is noisy but cheap (e.g. a computer program), in order to minimize total sampling expenditures. Our analysis gives necessary conditions for the Hybrid-Offset design and specifies optimal sample sizes for both annotators. Simulations on data from a coral reef survey program indicate that the Hybrid-Offset design outperforms several alternative sampling designs. In particular, sampling expenditures are reduced 50% compared to the Conventional design currently deployed by the coral ecologists.

1 Introduction

Using random sampling to estimate the mean of a population is a fundamentally important method to the sciences and society at large, and has been studied extensively [24, 4]. Deployment of any random sampling design requires collection of some number of observations sampled randomly from nature. In the ecological sciences, this was traditionally done *in situ* by an expert. Recently, advances in robotics, sensor technology, digital storage, and information technology have enabled rapid collection of samples in digital format, such as images [19, 10] or audio [11]. The popularity of digital sample collection can be attributed to three key factors: it creates a permanent record; it can be done cheaply using automated sampling vehicles or non-expert personnel; and it is generally fast. However, such samples (e.g. a photoquadrat of the forest floor) typically require annotation by an expert in order to reveal the desired quantity of interest (e.g. a count of insects). Such annotation work can be slow, tedious, expensive, and prone to error [13, 14].

Concurrent with the development of automated collection methods, advances in computer-vision and computer-audition have enabled automation of said annotation work. Such methods often rely on machine learning where expert annotated archived data sets are utilized to train automated annotators. Automation is a compelling low-cost alternative to expert annotations, but it's generally less reliable and may be biased [2, 8, 25]. This is particularly problematic if the probability density of the archived data differs from the density of the data to be sampled [17].

Crowdsourcing offers another low-cost alternative to expert annotation for e.g. document or image annotation [27, 15]. Crowdsourced annotations can be noisy, and much work has been devoted to improving the quality of such annotations. This is generally done either by carefully designing the tasks given to the crowd workers [15], or by collecting multiple crowd annotations for the same sample and then modeling, and compensating for, the annotation errors [27].

We consider the problem of estimating a population mean under these new cost-structures of data collection and annotation. This is formalized as follows: Given a procedure for collecting random samples, $x_i \in \mathcal{X}$ (e.g. images or audio recordings), each with an associated quantity of interest (value), $y_i \in \mathcal{Y} \subseteq [0, 1]$; and two annotators: a primary, $f_a(x_i) \in \mathcal{Y}$ which is accurate but costly (e.g. a human expert), and an auxiliary, $f_b(x_i) \in \mathcal{Y}$ which is cheap but noisy (e.g. a computer program, lay-person, or crowd worker). Our goal is to derive a sampling design that achieves unbiased estimates of the population mean ($E[y_i]$) at a target error and confidence, while minimizing total cost of collection and annotation. In particular, we investigate the optimal balance between the number of samples annotated by the primary and auxiliary annotators. This work is, to the best of our knowledge, the first to consider this problem.

A key challenge is to define a procedure that can correct for the potential bias of the auxiliary annotator. This is difficult as we cannot assume any prior knowledge of the underlying probability density from which the samples are drawn. Indeed, if this density was known, the population mean could be estimated directly, making the sampling work unnecessary. If the auxiliary annotator is based on machine learning and trained on archived data with a different probability density, the problem of transfer learning arises for which the generalization bounds of statistical learning theory generally do not apply [17]. Methods for bias-correction have been proposed independently by [23, 9], that do not require knowledge of the full underlying probability density (of the sampled data) but only of the conditional probabilities of a label given a sample. However, as we shall see, this information may not always be at hand.

The key contribution of this work is an analysis of a Hybrid-Offset design that directly models the offset (bias) of the auxiliary annotator. It is “Hybrid” because it requires a subset of the samples to be annotated by both annotators, and is unbiased if the samples are independent and identically distributed, which they are by construction under random sampling. As demonstrated by simulations on coral reef survey data, the Hybrid-Offset design is cost-effective and robust. In particular, it outperforms a Hybrid-Ratio design which utilizes the ratio-estimator commonly used in the sampling literature [24, 20]. It also outperforms several designs that only rely on one of the annotators, including the design currently used by the coral ecologists. We believe that the Hybrid-Offset can be widely utilized, in particular for ecological surveys relying on digital samples [29, 16, 19, 1, 11]. Other contributions of this work are: (1) an analysis of the bias-correction method of [23, 9] in the context of random sampling; and (2) an improved machine-learning method, based on convolutional neural networks, for annotation of coral reef survey images.

1.1 Related work

Our work is most closely related to the literature on survey sampling with auxiliary data [24, 20]. In that context, “auxiliary data” is typically not a direct estimate of the variable of interest but some other, related quantity. For example, if the variable of interest, y_i is the number of animals per plot, the auxiliary data can be the plot area, vegetation type, or plot elevation. In this literature, auxiliary data is typically incorporated using a ratio estimator which can reduce the estimation errors if there is an approximately linear relationship between the auxiliary data and the variable of interest [24]. However, the ratio estimator is design-biased, and analysis of the estimator variance typically assumes that auxiliary data is available for the whole population [24, 20]. In this work, in contrast, we do not assume that the auxiliary data is available for all samples, but that this data can be acquired at a cost by collecting additional samples. This additional cost is then taken into account when deriving optimal sampling sizes. Another key difference is that we do not use the ratio estimator but an offset estimator which directly estimates the bias of the auxiliary annotator. The offset estimator is design-unbiased and allows for a more straight-forward analysis.

Another line of related work utilizes stratified random sampling [3], importance sampling [21] or generative models of the classifier score distribution [28], to achieve cost-effective estimates of classifier performance on new data. The work of Garnett et al. [5] is particularly relevant, and investigates methods for active selection of samples in order to estimate class proportions. However, these methods all operate on a fixed set of samples. In contrast, we include the sample collection in our model which enables a joint minimization of annotation and collection costs.

Our work is also related to active learning and transfer learning. It is related, in particular, to recent work on active transfer learning where labels are queried to optimize classifier performance in a

target domain [26]. A key difference between that work and ours is that active learning methods optimize the labeling effort to create the best *classifier* (which then can, presumably, be used to label more data and in order to estimate the desired data-products). In contrast, we directly optimize the labeling effort to derive the desired *data-product* (i.e. the population mean).

2 Preliminaries

2.1 Problem Setup

We denote by $\mu_p \equiv E[y_i]$ and $\sigma_p^2 \equiv \text{var}[y_i]$ the first and second moments of the (unknown) probability density function of the values. The values are sampled randomly from nature, and are therefore i.i.d. We further denote by $f_a : \mathcal{X} \rightarrow \mathcal{Y}$ the primary, and by $f_b : \mathcal{X} \rightarrow \mathcal{Y}$ the auxiliary annotator, $\epsilon_{a,i} \equiv f_a(x_i) - y_i$ the error of f_a on sample i , $\mu_a \equiv E[\epsilon_{a,i}]$ and $\sigma_a^2 \equiv \text{var}[\epsilon_{a,i}]$. Similarly, $\mu_b \equiv E[\epsilon_{b,i}]$ and $\sigma_b^2 \equiv \text{var}[\epsilon_{b,i}]$. Note that we do *not* make any assumptions on the underlying probability densities of the sample values or annotator errors.

We denote by n_a and n_b the number of samples annotated by f_a and f_b , respectively. The number of collected samples is given by $\max(n_a, n_b)$ since samples needs to be annotated to provide any information, and conversely, needs to be collected in order to be annotated. We denote by c_c, c_a and c_b the cost per sample for collection, annotation by f_a and annotation by f_b , respectively. The ‘accurate and expensive’ characteristics of f_a are operationalized by letting $\sigma_a^2 < \sigma_b^2$ and $c_a > c_b$.

We can now precisely state our goal: Given costs c_c, c_a and c_b , and two annotators, f_a and f_b , derive a sampling design that estimates the population mean, μ_p by defining the number of annotated samples (n_a and n_b), so that $E[\hat{\mu}_p] = \mu_p$ and $\Pr(|\hat{\mu}_p - \mu_p| > d) < \delta$, for a target error, d and confidence, δ . The utility of the sampling design is evaluated by the Total Sampling Cost (TSC), b :

$$b(n_a, n_b) = c_a n_a + c_b n_b + \max(n_a, n_b) c_c. \quad (1)$$

We make three assumptions. First, we assume that the number of collected samples is small in comparison with the total size of the population which allows us to omit the finite-population correction factor [24]. Second, we assume that the primary annotator, f_a is unbiased, i.e. $\mu_a = 0$, and that the correlation between $\epsilon_{a,i}$ and y_i is negligible. Third, because the two annotators are independent entities, we assume zero correlation between the annotator errors $\epsilon_{a,i}$ and $\epsilon_{b,i}$. However, we do *not* make any assumptions on the correlation between the auxiliary annotation errors $\epsilon_{b,i}$ and the sample value y_i , which may be large. All proofs are in the Appendix.

2.2 Conventional design

We denote by ‘conventional’, a sampling design where all collected samples, x_i are annotated by the primary annotator f_a , i.e. $n_b = 0$. In such design, an unbiased estimator of μ_p is given by

$$\hat{\mu}_p = \frac{1}{n_a} \sum_{i=1}^{n_a} f_a(x_i), \quad (2)$$

with variance

$$\text{var}[\hat{\mu}_p] = \frac{1}{n_a} (\sigma_p^2 + \sigma_a^2). \quad (3)$$

The variance $(\sigma_p^2 + \sigma_a^2)$ is often unknown, and must be estimated by the sample variance of $f_a(x_1), \dots, f_a(x_{n_a})$. The sample size, n_a needs to be large enough to ensure that $\Pr(|\hat{\mu}_p - \mu_p| > d) < \delta$, for a target error d and confidence δ . From the Central Limit Theorem, this is satisfied when

$$\zeta_\delta \sqrt{\text{var}[\hat{\mu}_p]} \leq d, \quad (4)$$

where ζ_δ is the upper $1 - \delta/2$ point on the standard normal distribution curve [24]. The target sample size is given by inserting (3) into (4) yielding

$$n_a^* = \frac{\zeta_\delta^2}{d^2} (\sigma_p^2 + \sigma_a^2), \quad (5)$$

for a TSC: $(c_c + c_a) n_a^*$.

3 Hybrid-Offset design

Now consider a hybrid design where $n_b \geq n_a^*$ samples are collected and annotated by the auxiliary annotator f_b , and where a subset $n_a \leq n_b$ is also annotated by the primary annotator f_a . An offset estimator of μ_p under this design is given by

$$\hat{\mu}_p = \frac{1}{n_b} \sum_{i=1}^{n_b} f_b(x_i) - \hat{\mu}_b. \quad (6)$$

The offset estimator is unbiased and an unbiased estimate of μ_b is given by

$$\hat{\mu}_b = \frac{1}{n_a} \sum_{i=1}^{n_a} f_b(x_i) - f_a(x_i). \quad (7)$$

The variance of $\hat{\mu}_p$ is given by

$$\text{var}[\hat{\mu}_p] = \frac{1}{n_b}(\sigma_p^2 - \sigma_b^2) + \frac{1}{n_a}(\sigma_a^2 + \sigma_b^2), \quad (8)$$

and notably does *not* depend on the covariance between y_i and $\epsilon_{b,i}$. This follows directly from the derivations of (8), which are provided in Appendix A. We denote by “Hybrid-Offset” a design that balances n_a and n_b to minimize the TSC.

If the costs are such that a large number of samples, n_b can be collected and annotated by the auxiliary annotator, the magnitude of the first term in (8) becomes small and the sampling error depends mainly on the auxiliary annotation error, σ_b^2 . In contrast, the conventional sampling design depends mainly on the data variance, σ_p^2 (3). This is compelling because while the data variance is a fixed constant of nature, the auxiliary annotation errors depend on the choice and quality of the auxiliary annotator, which we control. It also leads to our first result:

Theorem 1. *For any fixed $n_b > n_a$, the variance of the Hybrid-Offset estimator (8) is smaller than the variance of the Conventional design (3) if and only if $\sigma_b^2 < \sigma_p^2$.*

Theorem 1 implies that the uncertainty introduced by f_b can be compensated for by using more samples if and only if $\sigma_b^2 < \sigma_p^2$. However, the additional collection of samples is only economical for certain cost functions, and should in the general case be determined by comparing the TSC of the two designs. To determine the TSC of the Hybrid-Offset design we begin by deriving optimal sample sizes n_b and n_a . By combining (8) and (4), and solving for equality, the following trade-off between n_b and n_a is derived

$$n_a = \frac{\sigma_b^2 + \sigma_a^2}{\frac{d^2}{\zeta_\delta^2} - \frac{1}{n_b}(\sigma_p^2 - \sigma_b^2)}. \quad (9)$$

Example trade-off curves demonstrate how the amount of primary annotations can be reduced by increasing the amount of auxiliary annotations (Fig. 1A). Note that if $n_b = n_a^*$, it follows from (9) that $n_a = n_a^*$, and the Hybrid-Offset design reduces to the Conventional design. The optimal operating point along the n_b, n_a trade-off curve can be derived by minimizing the TSC. Using (9) to eliminate n_a , the TSC becomes

$$b(n_a, n_b) = (c_c + c_b) \left(n_b + \frac{k(\sigma_b^2 + \sigma_a^2)}{\frac{d^2}{\zeta_\delta^2} - \frac{1}{n_b}(\sigma_p^2 - \sigma_b^2)} \right), \quad (10)$$

where $k = \frac{c_a}{c_c + c_b}$ is the relative cost of f_a . The optimal sample size, n_b is given by minimizing (10) under the constraint that $n_b \geq n_a^*$. This yields the following theorem:

Theorem 2. *Under a Hybrid-Offset design (6), the optimal auxiliary annotation size is*

$$n_b = \max \left[\frac{\zeta_\delta^2}{d^2} \left(\sigma_p^2 - \sigma_b^2 + \sqrt{k(\sigma_b^2 + \sigma_a^2)(\sigma_p^2 - \sigma_b^2)} \right), n_a^* \right] \quad (11)$$

The corresponding primary annotation size is given by (9).




Figure 1: (A) Amount of auxiliary (n_b), and primary (n_a) annotation required to achieve error $d \leq 0.02$ at 95% confidence for $\sigma_p = 0.2$, $\sigma_a = 0.02$, and $\sigma_b = \{0.25\sigma_p, 0.5\sigma_p, \sigma_p\}$ under the Hybrid-Offset sampling design. Solid gray line indicates $n_b = n_a$. Optimal operating points (11), for relative cost of primary annotation, $k = 10$ and $k = 100$ are marked with X on the $\sigma_b = 0.5\sigma_p$ curve. (B) Difference in Total Sampling Cost ($\text{TSC}_{\text{Conventional}} - \text{TSC}_{\text{offset}}$) for $\sigma_p = 0.2$, $\sigma_a = 0.02$, $\sigma_b = \{0.25\sigma_p, 0.5\sigma_p, 0.75\sigma_p\}$, $c_b = 0$ and $c_c = 1$, as a function of the relative cost of primary annotation, $k = c_a/(c_c + c_b)$. The threshold costs, σ_Δ are marked with black stumps. If $k < \sigma_\Delta$ the sampling designs and costs are identical.

Using the optimal sample sizes, the TSC of Hybrid-Offset sampling can be calculated from (10) and compared to the TSC of Conventional sampling in order to determine the most cost-effective design. For the important special case where $c_b = 0$ (which can occur e.g. if f_b is a computer algorithm) the following theorem applies:

Theorem 3. *If $c_b = 0$, the TSC of Hybrid-Offset sampling is smaller than the TSC of Conventional sampling if and only if $k > \sigma_\Delta$, where $k = \frac{c_a}{c_c + c_b}$ and $\sigma_\Delta = \frac{\sigma_a^2 + \sigma_b^2}{\sigma_p^2 - \sigma_b^2}$.*

For example, if the primary annotation errors are zero ($\sigma_a^2 = 0$), and the auxiliary errors are half as large as the data variance ($\sigma_b^2 = \sigma_p^2/2$), then $\sigma_\Delta = 1$. This means that Hybrid-Offset sampling is cheaper than Conventional sampling if $k > 1$, which occurs if the cost of collection is smaller than the cost of primary annotation. The difference in TSC between the two sampling designs is shown in Fig. 1B for various values of k , σ_p and σ_b .

Hybrid-Offset sample sizes can also be derived directly from a target TSC, b , and costs, c_a, c_b and c_c by minimizing (8) under the TSC constraint: $b = n_a c_a + n_b (c_b + c_c)$. This yields:

$$n_b = \frac{b \sqrt{c_a (c_b + c_c) \sigma_\Delta}}{c_a (c_b + c_c) \sigma_\Delta - (c_b + c_c)^2} - \frac{b}{c_a \sigma_\Delta - (c_b + c_c)}, \quad (12)$$

where as previous, $\sigma_\Delta = \frac{\sigma_a^2 + \sigma_b^2}{\sigma_p^2 - \sigma_b^2}$. The corresponding n_a is given by the TSC constraint.

4 Experiments

The proposed method is discussed in the context of an annual coral reef survey performed by the Moorea Coral Reef Long Term Ecological Research (MCR-LTER) program (<http://mcr.lternet.edu>). The program surveys six sites across the island of Moorea in French Polynesia. At each site, three habitats are surveyed: the fringing-reef and two habitats on the outer reef at 10 and 17 meter depth respectively, for a total of 18 sampling “units”. In each unit the goal, as dictated by the ecologists, is to estimate the percent cover of key benthic substrates, such as coral and algae. These data provide important information about the ecology when compared across sites, habitats and years. To estimate the percent cover for each unit, ecologists capture photographs (*in-situ* by a research-diver) at random locations along five line-transects at each site. The photographs are then annotated in order to estimate the percent cover for each photograph. This is done through random point sampling in which the substrate is identified at 200 random point locations in each photograph [18]. This procedure of using random sampling to annotate each collected random sample is commonly referred to as two-stage sampling [24, 4].

For the purpose of this discussion we will focus on the estimation quality of percent coral cover for each unit, and investigate the effect of using an Hybrid-Offset design in place of the Conventional

designs currently in use. For the simulations we use the Moorea Labeled Corals dataset, which is publicly available¹ and contains the full-resolution images and annotations from the LTER-MCR surveys conducted 2008 and 2009. We use the data from 2008 to train the auxiliary annotators, and estimate the sampling parameters, and the data from 2009 to run the sampling simulations.

As described above, each photograph, x_i , is a random sample with a corresponding coral cover $y_i \in [0, 1]$, and $n_a = 40$ in each unit. The coral cover estimated by the expert annotator, $f_a(x_i)$ is highly accurate [14], and we therefore use $\sigma_a^2 = 0$ in the simulations below. We do not account for approximation errors introduced by the second-stage (point-annotation) method, as this have been shown to have very limited effect on the final mean estimator [4, 18]. Manual annotation requires ~ 10 minutes per image to complete, while collection is quicker, with the 40 samples in a sampling-unit captured in a single 40 minute dive. With these parameters, the TSC for each unit is approximately $40(10 + 1)$ minutes, or 7.3 person-hours.

Auxiliary annotators: We use two auxiliary annotators, the “texton”-based classifier proposed by [2] which is publicly available, and a novel annotator based on convolutional neural networks (CNN) [7]. Both of these methods operate on $p \times p$ pixel images patches, and are denoted $f_{b^*}^{(\Psi)}(x_{i,j}) : \mathbb{R}^{p \times p \times 3} \rightarrow \{0, 1\}$, for $\Psi \in \{\text{texton}, \text{CNN}\}$, where $x_{i,j}$ is a patch from image i around random point-location j , and where output 1 indicates ‘coral’, and 0 ‘other’. The output of the auxiliary annotator, $f_b^{(\Psi)}(x_i)$ is given by averaging the point-classifications for the 200 random points in each image:

$$f_b^{(\Psi)}(x_i) = \frac{1}{200} \sum_{j=1}^{200} f_{b^*}^{(\Psi)}(x_{i,j}). \quad (13)$$

To develop the CNN based point-classifier, $f_b^{(\text{CNN})}$, we adopted a 16 layer CNN model developed for image classification [22]. The VGG16 model is publicly available² and operates on 226×226 RGB images. To fine-tune this model for coral classification, we cropped 226×226 patches from the 2008 images at each of the 200 annotated point locations. These cropped patches were used together with mirrored and rotated (by 90, 180 and 270 degrees) versions as training-data to fine-tune the weights of the VGG16 model largely following the procedure proposed by [7]. Classification was performed by cropping patches from the test images, and propagating them through the network.

Cost analysis: Using the data from 2008, the sample variance of the percent coral cover $\hat{\sigma}_p$ was 0.16 ± 0.1 (mean \pm SD, $n = 18$) across the units, meaning that for an average unit, a 95% confidence interval of mean coral cover, from the $n_a = 40$ samples, is 5.8%. By cross-validation on the data from 2008, the auxiliary annotator errors of $f_b^{(\text{texton})}$, $\hat{\sigma}_b$, was estimated as 0.047 ± 0.024 (mean \pm SD, $n = 18$). Since $\hat{\sigma}_p < \hat{\sigma}_b$ the Hybrid-Offset design is likely more cost-effective than a Conventional design (Theorem 1). Indeed, with $d = 5.8\%$, $\delta = 0.05$, the optimal Hybrid-Offset design is to collect 53 samples and manually annotate 5 (Theorem 2), for a TSC of 2.3 person-hours per unit; a 68.7% reduction compared to the Conventional design.

Simulation details: In order to validate the expected cost-savings, simulations were carried out on the coral survey images collected and annotated during 2009. Using the variance estimates from the 2008 data, sample sizes were determined for TSC, $b = 1, 2, \dots, 15$ person-hours using Eq. (12). For each of the 18 units, for each budget, and for 500 iterations, the required number of images were drawn randomly with replacement from the images pertaining to that unit³, and mean estimates $\hat{\mu}_p$ were calculated. From these estimates the mean error (bias): $\frac{1}{500} \sum_{k=1}^{500} \hat{\mu}_p - \mu_p^0$ and mean absolute error (MAE): $\frac{1}{500} \sum_{k=1}^{500} |\hat{\mu}_p - \mu_p^0|$ of each method was calculated by comparing to the “ground-truth” cover μ_p^0 , which was estimated from the expert annotations of the 40 images in the unit. In addition to the Conventional and Hybrid-Offset designs, we include three other designs, which are defined below. All sampling designs are evaluated using both $f_b^{(\text{texton})}$ and $f_b^{(\text{CNN})}$.

¹ <http://vision.ucsd.edu/content/moorea-labeled-corals>

² <https://github.com/BVLC/caffe/wiki/Model-Zoo>

³ Sampling with replacement was used to avoid finite-population artifacts due to the limited pool of 40 images per unit. In an actual application, there would be a large number of locations from where to capture the images and finite-population correction would not be needed.




Figure 2: **Simulations results on the coral reef survey data.** Results displayed as mean \pm SE for $n = 18$ sampling units and for TSC (budgets) between 1 and 15 person-hours per unit. Estimator bias is shown for the $f_b^{(\text{texton})}$ auxiliary annotator in panel (A), and for $f_b^{(\text{CNN})}$ in panel (B). Mean Absolute Error (MAE) is shown in panel (C), where estimates using $f_b^{(\text{texton})}$ are indicated with solid lines, and $f_b^{(\text{CNN})}$ with dotted. Gray dash-dotted lines indicate the current operating point of the Conventional sampling design currently utilized by the ecological monitoring program. The MAEs of the Bias-Corrected estimates were $> 3\%$ (shown in Fig. S1).

Hybrid-Ratio Design: Auxiliary information is commonly incorporated using a ratio-estimator which assumes a linear relation between auxiliary values and the values of interest [24, 20]. The mean estimate of a ratio-estimator is $\hat{\mu}_p = \frac{1}{n_b} \sum_{i=1}^{n_b} \hat{r} f_b(x_i)$, where the ratio, r , is estimated as

$$\hat{r} = \frac{\sum_{i=1}^{n_a} f_a(x_i)}{\sum_{i=1}^{n_a} f_b(x_i)} \quad (14)$$

The ratio-estimator is design-biased, and there are several approximations of the ratio-estimator variance [20]. We do not analysis the ratio-estimator, but include it in the simulations using a Hybrid-Ratio design with the same sample sizes as the Hybrid-Offset design (12). We set $\hat{r} = 1$ if the nominator or denominator of (14) is zero in order to avoid ill-defined estimates of r .

Auxiliary Design: Using only the auxiliary annotator, μ_p can be estimated as $\hat{\mu}_p = \frac{1}{n_b} \sum_{i=1}^{n_b} f_b(x_i)$, with sample sizes $n_a = 0, n_b = b/c_c$. This estimator is not design-unbiased, but will have small variance since a large number of samples can be collected and annotated cheaply.

Auxiliary Bias-Corrected Design: The expected value of $f_{b^*}(x)$ for a sample x with value $y \in \{0, 1\}$ is

$$E \begin{bmatrix} f_{b^*}(x) \\ 1 - f_{b^*}(x) \end{bmatrix} = \begin{bmatrix} \alpha & 1 - \beta \\ 1 - \alpha & \beta \end{bmatrix} \begin{bmatrix} y \\ 1 - y \end{bmatrix}, \quad (15)$$

where α and β are the classifier sensitivity and specificity respectively. As noted by [23, 9], an unbiased estimate of y is given by inverting the confusion-matrix in the center of (15) yielding: $\tilde{y} = \frac{f_{b^*}(x) + \beta - 1}{\alpha + \beta - 1}$. We denote this operation as “bias-correction”, since the corrected values are unbiased in expectation. Since f_b is a linear combination of f_{b^*} (13), μ_p can be estimated as $\hat{\mu}_p = \frac{1}{n_b} \sum_{i=1}^{n_b} \frac{f_b(x_i) + \beta - 1}{\alpha + \beta - 1}$. Using cross-validation on the data from 2008, specificity and sensitivity were estimated as $\hat{\alpha} = 0.738, \hat{\beta} = 0.963$ for $f_{b^*}^{(\text{texton})}$, and $\hat{\alpha} = 0.740, \hat{\beta} = 0.968$ for $f_{b^*}^{(\text{CNN})}$, and used for the bias-correction. An analysis of the variance of a mean estimator from abundance-corrected values is provided in Appendices E & F.

5 Results & Discussion

As expected, the Conventional and Hybrid-Offset designs were unbiased (Fig. 2A). The Hybrid-Ratio design also had low bias, except for smaller budgets, where, as mentioned above, the ratio-estimator (14) may be ill-posed. The Auxiliary Bias-Corrected design was, in-fact, biased (Fig. 2A,

B). This may seem surprising, but the corrected estimates are only unbiased if the sensitivity, α and specificity, β are known [23, 9]. These results thus indicate that $\hat{\alpha}$ and $\hat{\beta}$ estimated from the 2008 data were not valid for the sampling units from 2009. This may be due to domain-shifts, which can severely affect the performance of machine-learning based classifiers [25, 17]. One way to circumvent this problem is to use Bias-Correction in a Hybrid design, and estimate $\hat{\alpha}$ and $\hat{\beta}$ from the n_a samples annotated by both annotators. However, as shown in Appendix H, such Hybrid-Bias-Correction design is inferior to the Hybrid-Offset design. Finally, and less surprisingly, the un-corrected Auxiliary design was biased, although to a lesser extent for $f_b^{(\text{CNN})}$ (Fig. 2A, B).

The MAE of the Conventional design at the 7.3 person-hour budget currently utilized by the MCR-LTER program was $1.63 \pm .32\%$ (mean \pm SE, $n = 18$; Fig. 2B). This was outperformed by the Hybrid-Offset and Hybrid-Ratio designs, which utilized both annotators. At the 7.3 person-hour budget, the Hybrid-Offset estimator MAE was $1.08 \pm .16\%$ when relying on $f_b^{(\text{texton})}$ and $1.00 \pm .18\%$ when relying on $f_b^{(\text{CNN})}$, which is significantly lower than for the Conventional design. Conversely, the MAE that the Conventional design achieved at 7.3 person-hours, can be achieved by the Hybrid-Offset design at around 3.5 person-hours; a $\sim 50\%$ cost reduction. The Hybrid-Ratio design, while comfortably better than the Conventional design, performed worse than the Hybrid-Offset design for all budgets. This may be because the sample sizes were optimal for the offset-estimator and not the ratio-estimator. The Auxiliary and Auxiliary Bias-Corrected designs, which both relied only on the auxiliary annotator, performed weaker than the Hybrid and Conventional designs. The MAE of the Bias-Corrected design was $> 3\%$ for all budgets, indicating that the correction method of [23, 9] was ineffective. The Auxiliary design performed poorly when relying on $f_b^{(\text{texton})}$ (MAE $> 2\%$), but better when relying on $f_b^{(\text{CNN})}$, barely outperforming the Conventional design at the 7.3 person-hour budget (Fig. 2C).

It is also clear from the simulations that $f_b^{(\text{CNN})}$ outperformed $f_b^{(\text{texton})}$, reducing the errors for the Hybrid and Auxiliary estimators. This is expected as CCN-based methods have recently achieved state of the art performance on several visual recognition tasks [7]. As new, and stronger classification methods are developed, the requirement of Theorem 1 will be satisfied for an increasing number of applications, suggesting increasing utility of the offset sampling design.

We have used linear cost-functions throughout this work, with a fixed cost per sample. In reality, the cost per sample is likely to decrease when more samples are collected. This is true, in particular for spatial surveys where the sample collector will, on average, have shorter distance to travel between the samples. Since Hybrid sampling designs require larger number of collected samples, the cost-savings estimated by our simulations should be considered a lower-bound on the actual cost-savings.

6 Conclusion

We have investigated the implications of modeling and incorporating the cost and accuracy of two annotators in random sampling designs for population mean estimation, and shown that significant cost-savings are possible using data from a coral reef survey. However, the derived formulations are general and applicable in many other situations. These include other marine surveys of e.g. fish [29] or plankton [16], terrestrial surveys of crops [10], forests [6], rangelands [19] and deserts [1]; and audio-based surveys of e.g. marine mammal or bird populations [11].

To the best of our knowledge, this is the first work that models automated annotation as part of the sampling design and we believe that there are several interesting directions of future work, notably with respect to stratified and sequential sampling procedures.

References

- [1] S. Archer, C. Scifres, C. Bassham, and R. Maggio. Autogenic succession in a subtropical savanna: conversion of grassland to thorn woodland. *Ecol Monograph*, 1988.
- [2] O. Beijbom, P. J. Edmunds, D. I. Kline, B. G. Mitchell, and D. Kriegman. Automated annotation of coral reef survey images. In *CVPR*, 2012.
- [3] P. N. Bennett and V. R. Carvalho. Online stratified sampling: evaluating classifiers at web-scale. In *CIKM*, 2010.

- [4] W. E. Deming. *Some theory of sampling*. Courier Dover Publications, 1966.
- [5] R. Garnett, Y. Krishnamurthy, X. Xiong, J. Schneider, and R. Mann. Bayesian optimal active search and surveying. *ICML*, 2012.
- [6] S. Getzin, K. Wiegand, and I. Schöning. Assessing biodiversity in forests using very high-resolution images and unmanned aerial vehicles. *Methods Ecol. Evol.*, 2012.
- [7] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [8] D. J. Hand. Classifier technology and the illusion of progress. *Statistical Science*, 2006.
- [9] D. J. Hopkins and G. King. A method of automated nonparametric content analysis for social science. *Am J Polit Sci*, 2010.
- [10] E. R. Hunt, W. D. Hively, S. J. Fujikawa, D. S. Linden, C. S. Daughtry, and G. W. McCarty. Acquisition of nir-green-blue digital photographs from unmanned aircraft for crop monitoring. *Remote Sens*, 2010.
- [11] M. P. Johnson and P. L. Tyack. A digital acoustic recording tag for measuring the response of wild marine mammals to sound. *J. Ocean. Eng.*, 2003.
- [12] K. E. Kohler and S. M. Gill. Coral point count with excel extensions (cpce): A visual basic program for the determination of coral and substrate coverage using random point count methodology. *Comput Geosci*, 2006.
- [13] N. MacLeod, M. Benfield, and P. Culverhouse. Time to automate identification. *Nature*, 2010.
- [14] R. Ninio, J. Delean, K. Osborne, and H. Sweatman. Estimating cover of benthic organisms from underwater video images: variability associated with multiple observers. *MEPS*, 2003.
- [15] J. Noronha, E. Hysen, H. Zhang, and K. Z. Gajos. Platemate: crowdsourcing nutritional analysis from food photographs. In *UIST*, 2011.
- [16] R. J. Olson and H. M. Sosik. A submersible imaging-in-flow instrument to analyze nano-and microplankton: Imaging flowcytobot. *Limnol. Oceanogr. Methods*, 2007.
- [17] S. J. Pan and Q. Yang. A survey on transfer learning. *KDE*, 2010.
- [18] E. Pante and P. Dustan. Getting to the point: Accuracy of point count in monitoring ecosystem change. *J Mar Bio*, 2012.
- [19] A. Rango, A. Laliberte, J. E. Herrick, C. Winters, K. Havstad, C. Steele, and D. Browning. Unmanned aerial vehicle-based remote sensing for rangeland assessment, monitoring, and management. *J Appl Remote Sens*, 2009.
- [20] R. M. Royall and W. G. Cumberland. An empirical study of the ratio estimator and estimators of its variance. *J Am Stat Assoc*, 1981.
- [21] C. Sawade, N. Landwehr, S. Bickel, and T. Scheffer. Active risk estimation. In *ICML*, 2010.
- [22] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv*, 2014.
- [23] A. Solow, C. Davis, and Q. Hu. Estimating the taxonomic composition of a sample when individuals are classified with error. *Mar Ecol Prog Ser*, 2001.
- [24] S. K. Thompson. *Sampling*. John Wiley & Sons, Inc., 2012.
- [25] A. Torralba and A. A. Efros. Unbiased look at dataset bias. In *CVPR*, 2011.
- [26] X. Wang, T.-K. Huang, and J. Schneider. Active transfer learning under model shift. In *ICML*, 2014.
- [27] P. Welinder, S. Branson, P. Perona, and S. J. Belongie. The multidimensional wisdom of crowds. In *NIPS*, 2010.
- [28] P. Welinder, M. Welling, and P. Perona. A lazy man’s approach to benchmarking: Semisupervised classifier evaluation and recalibration. In *CVPR*, 2013.
- [29] T. J. Willis and R. C. Babcock. A baited underwater video system for the determination of relative density of carnivorous reef fish. *Mar Freshwater Res*, 2000.

Appendices: Random Sampling in an Age of Automation: Minimizing Expenditures through Balanced Collection and Annotation

A Variance of the offset estimator

To derive the variance of $\hat{\mu}_p$ we begin by expanding out and separating all terms,

$$\hat{\mu}_p = \frac{1}{n_b} \sum_{i=1}^{n_b} f_b(x_i) - \hat{\mu}_b. \quad (\text{S1})$$

$$= \frac{1}{n_b} \sum_{i=1}^{n_b} f_b(x_i) - \frac{1}{n_a} \sum_{i=1}^{n_a} f_b(x_i) - f_a(x_i) \quad (\text{S2})$$

$$= \frac{1}{n_b} \sum_{i=1}^{n_b} y_i + \frac{n_a - n_b}{n_a n_b} \sum_{i=1}^{n_a} \epsilon_{b,i} + \frac{1}{n_b} \sum_{i=n_a+1}^{n_b} \epsilon_{b,i} + \frac{1}{n_a} \sum_{i=1}^{n_a} \epsilon_{a,i}, \quad (\text{S3})$$

after which the variance of $\hat{\mu}_p$ is given by:

$$\text{var}[\hat{\mu}_p] = \frac{1}{n_b} (\sigma_p^2 - \sigma_b^2) + \frac{1}{n_a} (\sigma_a^2 + \sigma_b^2) + \frac{2}{n_b} \sigma_{a,p} + 2 \frac{n_a - n_b}{n_a n_b} \sigma_{a,b}. \quad (\text{S4})$$

The first two terms are the variance of the data and annotator errors. The third term is the covariance between the primary annotator errors and the sample values $\sigma_{a,p}$. Since, by assumption, the primary annotator is “accurate”, we can expect this term to be small, and can be omitted. The last term accounts for the covariance between the annotator errors $\sigma_{a,b}$. Since the two annotators are operating independently, this is assumed to be zero, and it can also be omitted. Interestingly, the third covariance term, $\sigma_{p,b}$, between the auxiliary annotator and the actual values, which may be significant, cancels out and does not affect the final expression.

B Proof of Theorem 1

Proof. It follows from (3) and (8) that if $\sigma_b = \sigma_p \Rightarrow \text{var}[\hat{\mu}_p] = \text{var}[\hat{\mu}_p]$. Since $\text{var}[\hat{\mu}_p]$ increases monotonically with increasing σ_b (since $n_b > n_a$) the theorem follows trivially. \square

C Proof of Theorem 2

We start with the following lemma:

Lemma 1. *The cost function of offset sampling (10) is convex for $\sigma_b^2 < \sigma_p^2$, $n_b \geq n_a^*$ where n_a^* is given by (5).*

Proof. We begin by deriving the first and second derivatives of (10):

$$\frac{\partial t(n_a, n_b)}{\partial n_b} = (c_c + c_b) \left(1 - \frac{k(\sigma_b^2 + \sigma_a^2)(\sigma_p^2 - \sigma_b^2)}{(n_b \frac{d^2}{\zeta_\delta^2} - (\sigma_p^2 - \sigma_b^2))^2} \right) \quad (\text{S5})$$

$$\frac{\partial^2 t(n_a, n_b)}{\partial n_b^2} = (c_c + c_b) \left(\frac{2k(\sigma_b^2 + \sigma_a^2)(\sigma_p^2 - \sigma_b^2)(\frac{d^2}{\zeta_\delta^2})}{(n_b \frac{d^2}{\zeta_\delta^2} - (\sigma_p^2 - \sigma_b^2))^3} \right). \quad (\text{S6})$$

Since by assumption $\sigma_b^2 < \sigma_p^2$, the second derivative is positive for

$$n_b \geq n_a^* > \frac{\zeta_\delta^2}{d^2} (\sigma_p^2 - \sigma_b^2), \quad (\text{S7})$$

where the strict inequality requires either σ_a^2 or σ_b^2 to be non-zero. This concludes the proof. \square

Since, according to Theorem 1, offset sampling should only be considered if $\sigma_b^2 < \sigma_p^2$, and since by design, $n_b \geq n_a^*$, (10) is convex, and Theorem 2 follows by setting the first derivative to zero and solving for n_b .

D Proof of Theorem 3

Proof. The TSC of conventional sampling, under the assumption that $c_b = 0$, is given by $n_a(c_c + c_a)$ and the TSC of offset sampling is given by $n_b c_c + n_a c_a$. As noted previously, if $n_b = n_a = n_a^*$, offset sampling reduces to conventional sampling and the TSCs are equal. Since n_b minimize the offset TSC, which is convex for $n_b \geq n_a^*$ (Appendix C), it follows that the TSC of offset sampling is smaller than the TSC of conventional sampling if and only if $n_b > n_a^*$. The threshold, σ_Δ for when this occurs can be calculated by equating the two arguments inside the max operator of (11) and solving for k :

$$\frac{\zeta_\delta^2}{d^2}(\sigma_p^2 + \sigma_a^2) = \frac{\zeta_\delta^2}{d^2} \left(\sigma_p^2 - \sigma_b^2 + \sqrt{\sigma_\Delta(\sigma_b^2 + \sigma_a^2)(\sigma_p^2 - \sigma_b^2)} \right) \Rightarrow \quad (\text{S8})$$

$$\sigma_a^2 + \sigma_b^2 = \sqrt{\sigma_\Delta(\sigma_b^2 + \sigma_a^2)(\sigma_p^2 - \sigma_b^2)} \Rightarrow \quad (\text{S9})$$

$$\sigma_\Delta = \frac{\sigma_a^2 + \sigma_b^2}{\sigma_p^2 - \sigma_b^2}. \quad (\text{S10})$$

□

E Random Sampling with Abundance Correction

A third sampling design can be defined under one critical additional assumption. This assumption, which we will denote the ‘population drift assumption’, is that the performance of f_b can be defined in terms of a matrix of confusion, Q (which excludes real-valued output spaces \mathcal{Y}), and that Q is known *a priori* for the data to be sampled.

For these derivations we will let $\mathcal{Y} = \{0, 1\}$, meaning that the samples y are drawn from $\text{Ber}(\mu_p)$, a Bernoulli distribution with mean μ_p . This corresponds to annotating each sample x_i as containing or not containing the quantity of interest. In Appendix F we derive the statistics of auxiliary sampling in a two-stage sampling design [24, 4], where each sample x_i is annotated by second stage sampling, from which the corresponding y_i is obtained.

A matrix of confusion, Q characterizes the misclassification rates of f_b . In the binary case, Q is a two by two matrix

$$Q = \begin{bmatrix} \alpha & 1 - \beta \\ 1 - \alpha & \beta \end{bmatrix},$$

where α is the sensitivity and β the specificity. As noted independently by [23, 9] Q can be used to create an unbiased estimate of y , and we begin by recalling this procedure. The expected value of the auxiliary annotation f_b of a sample x with value y is

$$E \begin{bmatrix} f_b(x) \\ 1 - f_b(x) \end{bmatrix} = Q \begin{bmatrix} y \\ 1 - y \end{bmatrix},$$

and an unbiased estimator of y is given by inverting Q

$$\tilde{y} = \frac{f_b(x) + \beta - 1}{\alpha + \beta - 1}. \quad (\text{S11})$$

We refer to this as the ‘abundance corrected’ value, and derive a sampling procedure based on this correction. The variance of \tilde{y}_i , given the true value y_i is

$$\text{var}[\tilde{y}_i | y_i] = \frac{\text{var}[f_b(x_i) | y_i]}{(\alpha + \beta - 1)^2}, \quad (\text{S12})$$

which follows directly from (S11). We also note that

$$\text{var}[f_b(x_i) | y_i] = y_i \alpha (1 - \alpha) + (1 - y_i) (1 - \beta) \beta, \quad (\text{S13})$$

since if $y_i = 1$, $f_b(x_i) \sim \text{Ber}(\alpha)$, and if $y_i = 0$, $f_b(x_i) \sim \text{Ber}(1 - \beta)$. Combining (S12) and (S13), yields

$$\text{var}[\tilde{y}_i | y_i] = \frac{y_i \alpha (1 - \alpha) + (1 - y_i) (1 - \beta) \beta}{(\alpha + \beta - 1)^2}. \quad (\text{S14})$$

Finally, $\text{var}[\tilde{y}_i]$ is given by the law of total variance

$$\text{var}[\tilde{y}_i] = E[\text{var}(\tilde{y}_i|y_i)] + \text{var}[E(\tilde{y}_i|y_i)] \quad (\text{S15})$$

$$= \sigma_s^2 + \sigma_p^2, \quad (\text{S16})$$

where σ_p^2 is the data variance and

$$\sigma_s^2 = \frac{\mu_p \alpha (1 - \alpha) + (1 - \mu_p)(1 - \beta)\beta}{(\alpha + \beta - 1)^2}, \quad (\text{S17})$$

the variance introduced by the abundance correction. If the classifier is balanced, i.e. $\alpha = \beta$, (S17) simplifies to $\sigma_s^2 = \frac{\alpha(1-\alpha)}{(2\alpha-1)^2}$. Since \tilde{y}_i is an unbiased estimator of y_i , we can achieve an unbiased estimation of μ_p as

$$\hat{\mu}_p = \frac{1}{n_b} \sum_{i=1}^{n_b} \tilde{y}_i, \quad (\text{S18})$$

with variance, assuming that σ_s^2 and σ_p^2 are uncorrelated

$$\text{var}[\hat{\mu}_p] = \frac{1}{n_b} (\sigma_s^2 + \sigma_p^2), \quad (\text{S19})$$

and sample size

$$n_b = \frac{\zeta_\delta^2}{d^2} (\sigma_s^2 + \sigma_p^2). \quad (\text{S20})$$

Finally, the TSC, since $n_a = 0$, is given by

$$t(n_a, n_b) = (c_c + c_b)n_b. \quad (\text{S21})$$

The auxiliary sampling design requires annotation of n_b samples by the auxiliary annotator, but as it does not require any annotations by the primary annotator, the TSC can be low. The following theorem is given directly from (S21) and the cost function for conventional sampling.

Theorem 4. *For binary output spaces, $\mathcal{Y} = \{0, 1\}$, the TSC of auxiliary sampling is smaller than the TSC of conventional sampling if and only if*

$$k' > \frac{\sigma_p^2 + \sigma_s^2}{\sigma_p^2 + \sigma_a^2} \quad (\text{S22})$$

where $k' = (c_c + c_a)/(c_c + c_b)$.

If f_b is accurate then σ_s^2 is small and auxiliary sampling is cheaper than conventional sampling even for low primary annotation costs. For example, if $\alpha = \beta = 0.9 \Rightarrow \sigma_s^2 \approx 0.14$, $\sigma_p^2 = 0.04$, and $\sigma_a^2 = 0$, it suffices that k' is larger than 4.5, which is satisfied e.g. if $c_b = 0$ and $c_a > 3.5c_c$. If, on the other hand, $\alpha = \beta = 0.7 \Rightarrow \sigma_s^2 \approx 1.3$, k' must be larger than 33.5.

F Two-Stage Random Sampling with Abundance Correction

In two-stage sampling designs each first stage sample, x_i is again sampled randomly using some number, s of second-stage samples [24]. An analysis of the errors using such designs is provided by Deming [4]. Second stage sampling is commonly used e.g. in benthic surveys where each collected photoquadrat is annotated using random point sampling [12]. This protocol requires s points to be overlaid on each image at locations selected randomly with replacement. The substrate under each point is then annotated by an expert as pertaining to one of some number of classes. An unbiased estimator of the abundance (benthic cover) of each class for a certain sample can be derived by counting how many of the s annotations that were annotated as that class.

We derive the statistics of two-stage sampling under the population drift assumption, namely that each decision is made by some noisy annotator f_b , with known matrix of confusion. We will denote by x_{i1}, \dots, x_{is} the s locations to be annotated in each sample x_i , and $u_{ij} \in \{0, 1\}$ the true value associated with each location. The value of each first stage sample is approximated by $y_i = \sum_{j=1}^s u_{ij}$.

We do not make any assumptions on the probability density from which the first stage samples are drawn, but as previously let μ_p denote the expected value and σ_p^2 the variance. Given a classifier f_b with known matrix of confusion, an unbiased estimator of u_{ij} is given as previously by

$$\tilde{u}_{ij} = \frac{f_b(x_{ij}) + \beta - 1}{\alpha + \beta - 1}. \quad (\text{S23})$$

From this an unbiased estimator of \tilde{y}_i is given by

$$\tilde{y}_i = \frac{1}{s} \sum_{j=1}^s \tilde{u}_{ij}. \quad (\text{S24})$$

We have derived the variance of \tilde{y}_i for the special case where $s = 1$ in the main paper. Next, we show how to derive the variance of \tilde{y}_i for a general s by applying the law of total variation twice. We begin by noting that

$$\text{var}(\tilde{y}_i) = E[\text{var}(\tilde{y}_i|y_i)] + \text{var}[E(\tilde{y}_i|y_i)]. \quad (\text{S25})$$

The second term is simply $\text{var}[E(\tilde{y}_i|y_i)] = \text{var}[y_i] = \sigma_p^2$, and the first term can be expressed in terms of \tilde{u}_{ij}

$$\text{var}(\tilde{y}_i|y_i) = \frac{1}{s^2} \sum_{i=1}^s \text{var}(\tilde{u}_{ij}|y_i), \quad (\text{S26})$$

which can be expressed, by again using the law of total variation, as

$$\text{var}(\tilde{u}_{ij}|y_i) = E[\text{var}(\tilde{u}_{ij}|u_{ij}, y_i)] + \text{var}[E(\tilde{u}_{ij}|u_{ij}, y_i)]. \quad (\text{S27})$$

The second term of (S27) is simply given by $\text{var}[E(\tilde{u}_{ij}|u_{ij}, y_i)] = \text{var}[u_{ij}|y_i] = y_i(1 - y_i)$, but the first term is less obvious. Following Solow et al. [23], we first note that

$$\text{var}(f_b(x_{ij})|u_{ij}, y_i) = u_{ij}\alpha(1 - \alpha) + (1 - u_{ij})\beta(1 - \beta), \quad (\text{S28})$$

since if $u_{ij} = 1$, $f_b(x_{ij}) \sim \text{Ber}(\alpha)$, and if $u_{ij} = 0$, $f_b(x_{ij}) \sim \text{Ber}(1 - \beta)$. We then note that

$$\text{var}[\tilde{u}_{ij}|u_{ij}, y_i] = \frac{\text{var}[f_b(x_{ij})|u_{ij}, y_i]}{(\alpha + \beta - 1)^2}, \quad (\text{S29})$$

which follows directly from (S23), and also that $E[u_{ij}] = y_i$. Putting this together yields the following expression for the first term of (S27):

$$E[\text{var}(\tilde{u}_{ij}|u_{ij}, y_i)] = \frac{y_i\alpha(1 - \alpha) + (1 - y_i)\beta(1 - \beta)}{(\alpha + \beta - 1)^2}. \quad (\text{S30})$$

Putting this all together yields

$$\text{var}(\tilde{y}_i) = E[\text{var}(\tilde{y}_i|y_i)] + \text{var}[E(\tilde{y}_i|y_i)] \quad (\text{S31})$$

$$= E\left[\frac{1}{s^2} \sum_{i=1}^s \text{var}(\tilde{u}_{ij}|y_i)\right] + \sigma_p^2 \quad (\text{S32})$$

$$= \frac{1}{s^2} \sum_{i=1}^s E[\text{var}(\tilde{u}_{ij}|y_i)] + \sigma_p^2 \quad (\text{S33})$$

$$= \frac{1}{s^2} \sum_{i=1}^s E\left[\frac{y_i\alpha(1 - \alpha) + (1 - y_i)\beta(1 - \beta)}{(\alpha + \beta - 1)^2} + y_i(1 - y_i)\right] + \sigma_p^2 \quad (\text{S34})$$

$$= \frac{1}{s^2} \sum_{i=1}^s \left[\frac{\mu_p\alpha(1 - \alpha) + (1 - \mu_p)\beta(1 - \beta)}{(\alpha + \beta - 1)^2} + \mu_p(1 - \mu_p)\right] + \sigma_p^2 \quad (\text{S35})$$

$$= \frac{1}{s}[\sigma_s^2 + \mu_p(1 - \mu_p)] + \sigma_p^2, \quad (\text{S36})$$

where σ_s^2 is given by (S17). Interestingly, the variance of \tilde{y} approach σ_p for large number of secondary stage samples, s . This is to be expected under the assumption that f_b is perfectly modeled by

a known matrix of confusion Q . Since \tilde{y}_i is an average across s decisions, the variance introduced by the abundance correction cancels out with large values of s .

Finally, the total variance of $\hat{\mu}_p$ is given by

$$\text{var}[\hat{\mu}_p] = \frac{1}{n_b} \left(\frac{1}{s} [\sigma_s^2 + \mu_p(1 - \mu_p)] + \sigma_p^2 \right), \quad (\text{S37})$$

and the sample size by

$$n_b = \frac{\zeta_\delta^2}{d^2} \left[\frac{1}{s} [\sigma_s^2 + \mu_p(1 - \mu_p)] + \sigma_p^2 \right]. \quad (\text{S38})$$

G Supplementary results

Detailed simulation results are shown in Fig. S1.

H Bias-Correction Sampling With Unknown Confusion Matrix

The Auxiliary Bias-Correction design evaluated in the simulations assumes that the specificity and sensitivity of f_b is known *a priori* for the data to be sampled. This assumption is strong, and may not always hold. In such cases, one could rely on a Hybrid sampling design and use n_a samples annotated by both annotators to estimate $\hat{\alpha}$ and $\hat{\beta}$. In such Hybrid-Bias-Correction design, α and β can be estimated as

$$\hat{\alpha} = \frac{\sum_{i=1}^{n_a} f_{b^*}(x_i) f_{a^*}(x_i)}{\sum_{i=1}^{n_a} f_{a^*}(x_i)} \quad (\text{S39})$$

$$\hat{\beta} = \frac{\sum_{i=1}^{n_a} (1 - f_{b^*}(x_i))(1 - f_{a^*}(x_i))}{\sum_{i=1}^{n_a} (1 - f_{a^*}(x_i))}, \quad (\text{S40})$$

and an estimator of μ_p can be defined as

$$\hat{\mu}_p = \frac{1}{n_b} \left[\sum_{i=1}^{n_a} f_a(x_i) + \sum_{i=n_a+1}^{n_b} \frac{f_b(x_i) + \hat{\beta} - 1}{\hat{\alpha} + \hat{\beta} - 1} \right]. \quad (\text{S41})$$

However, we argue that such design is inferior to hybrid sampling for several reasons. First, the bias-corrected mean estimate of (S41) is biased if estimates of α and β are used in place of the true values [23]. Second, it is difficult to derive an analytical expression for $\text{var}[\hat{\mu}_p]$ that accounts for the variances of $\hat{\alpha}$ and $\hat{\beta}$. Without this expression, one cannot derive optimal sample sizes. Third, simulations detailed below indicate that the Hybrid-Offset design achieved lower errors for the same TSC for a wide array of parameters $(\alpha, \beta, \mu_p, n_a, n_b)$.

Simulations: For all combinations of $\alpha = [0.6, 0.8, 0.95]$, $\beta = [0.6, 0.8, 0.95]$, $\mu_p = [0.5, 0.75, 0.9]$, $n_a = [100, 150, \dots, 500]$, and $n_b = 1000$, the following simulation was performed. First, n_a samples $f_a(x_1), \dots, f_a(x_{n_a})$ were drawn from a Bernoulli (Ber) distribution with mean μ_p . For each $f_a(x_i) = 1$, $f_b(x_i)$ was drawn from $\text{Ber}(\alpha)$, and for each $f_a(x_i) = 0$, $f_b(x_i)$ was drawn from $\text{Ber}(1 - \beta)$. The parameters α, β and μ_b were then estimated according to (S39), (S40), and (7). Finally $n_b - n_a$ new samples $f_b(x_i)$ were drawn using the same procedure and used to estimate μ_p from (S41) and from (6). This procedure was repeated 2000 times and sample standard deviations were calculated. The signed difference between the standard deviations were calculated for each value of α, β, μ_p , and n_a , is shown in Fig. S2. These results indicate that the Hybrid-Offset design is more accurate than the Hybrid-Bias-Corrected design for all parameters.




Figure S1: **Supplementary Simulations Restuls** Results displayed as mean \pm SE for $n = 18$ sampling units, for TSC (budgets) between 1 and 15 person-hours per unit and for (Top) Mean Absolute Error and (Bottom) Mean Square Error




Figure S2: Difference in estimated standard errors between the Hybrid-Offset design and the Hybrid-Bias-Corrected design, for $n_b = 1000$, and different values of n_a , α , β , and μ_p . Note that all differences are negative, indicating that the sampling errors of the Hybrid-Offset design are smaller